

Sparse Estimation From Noisy Observations of an Overdetermined Linear System

Liang Dai^a and Kristiaan Pelckmans^a

^a*Division of Systems and Control, Department of Information Technology,
Uppsala University, Sweden
e-mail: liang.dai@it.uu.se, kristiaan.pelckmans@it.uu.se.*

Abstract

This note studies a method for the estimation of a finite number of unknown parameters from linear equations, which are perturbed by Gaussian noise. In case the unknown parameters have only few nonzero entries, the proposed estimator performs more efficiently than a traditional approach. The method consists of three steps: (1) a classical Least Squares Estimate (LSE), (2) the support is recovered through a Linear Programming (LP) optimization problem which can be computed using a soft-thresholding step, (3) a de-biasing step using a LSE on the estimated support set. The main contribution of this note is a formal derivation of an associated ORACLE property of the final estimate. That is, with probability 1, the estimate equals the LSE based on the support of the *true* parameters when the number of observations goes to infinity.

Key words: System identification; Parameter estimation; Sparse estimation.

1 Problem settings

This note considers the estimation of a sparse parameter vector from noisy observations of a linear system. The formal definition and assumptions of the problem are given as follows. Let $n > 0$ be a fixed number, denoting the dimension of the underlying parameter vector, and let $N > 0$ denote the number of equations ('observations'). The observed signal $\mathbf{y} \in \mathbb{R}^N$ obeys the following system:

$$\mathbf{y} = \mathbf{A}\mathbf{x}^0 + \mathbf{v}, \quad (1)$$

where the elements of the vector $\mathbf{x}^0 \in \mathbb{R}^n$ are considered to be the fixed but unknown parameters of the system. Moreover, it is assumed that \mathbf{x}_0 is s -sparse (i.e. there are s nonzero elements in the vector). Let $\mathcal{T} \subset \{1, \dots, n\}$ denote the support set of \mathbf{x}^0 (i.e. $\mathbf{x}_i^0 = 0 \Leftrightarrow i \notin \mathcal{T}$) and \mathcal{T}^c be the complement of \mathcal{T} , i.e. $\mathcal{T} \cup \mathcal{T}^c = \{1, 2, \dots, n\}$ and $\mathcal{T} \cap \mathcal{T}^c = \emptyset$. The elements of the vector $\mathbf{v} \in \mathbb{R}^N$ are assumed to follow the following distribution

$$\mathbf{v} \sim \mathcal{N}(0, cI_N), \quad (2)$$

where $0 < c \in \mathbb{R}$.

Applications of such setup appear in many places, to name a few, see the applications discussed in

Kump, Bai, Chan, Eichinger, and Li (2012) on the detection of nuclear material, and in Kukreja (2009) on model selection for aircraft test modeling (see also the Experiment 2 in Rojas and Hjalmarsson (2011) on the model selection for the AR model). In the experiment section, we will demonstrate an example which finds application in line spectral estimation, see Stoica and Moses (1997).

The matrix $\mathbf{A} \in \mathbb{R}^{N \times n}$ with $N > n$ is the sensing matrix. Such a setting (\mathbf{A} is a 'tall' matrix) makes it different from the setting studied in compressive sensing, where the sensing matrix is 'fat', i.e. $N \ll n$. For an introduction to the compressive sensing theory, see e.g. Donoho (2006); Candés and Wakin (2008).

Denote the Singular Value Decomposition (SVD) of matrix $\mathbf{A} \in \mathbb{R}^{N \times n}$ as

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (3)$$

in which $\mathbf{U} \in \mathbb{R}^{N \times N}$ satisfies $\mathbf{U}^T\mathbf{U} = I_N$, $\mathbf{V} \in \mathbb{R}^{n \times n}$ satisfies $\mathbf{V}^T\mathbf{V} = I_n$, and $\mathbf{\Sigma} \in \mathbb{R}^{N \times n}$ is a diagonal matrix $\mathbf{\Sigma} = \text{diag}(\sigma_1(\mathbf{A}), \sigma_2(\mathbf{A}), \dots, \sigma_n(\mathbf{A}))$. The results below make the following assumptions on \mathbf{A} :

Definition 1 *We say that $\{\mathbf{A} \in \mathbb{R}^{N \times n}\}_N$ are sufficiently rich if there exists a finite N_0 and $0 < c_1 \leq c_2$*

such that for all $N > N_0$ the corresponding matrices $\mathbf{A} \in \mathbb{R}^{N \times n}$ obey

$$c_1 \sqrt{N} \leq \sigma_1(\mathbf{A}) \leq \sigma_2(\mathbf{A}) \leq \dots \leq \sigma_n(\mathbf{A}) \leq c_2 \sqrt{N}, \quad (4)$$

where $\sigma_i(\mathbf{A})$ denotes the i -th singular value of the matrix \mathbf{A} , $c_1, c_2 \in \mathbb{R}^+$.

Note that the dependence of \mathbf{A} on N is not stated explicitly in order to avoid notational overload.

In Rojas and Hjalmarsson (2011) and Zou (2006), the authors make the assumption on \mathbf{A} that the sample covariance matrix $\frac{1}{N} \mathbf{A}^T \mathbf{A}$ converges to a finite, positive-definite matrix:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{A}^T \mathbf{A} = \mathbf{D} \succ 0. \quad (5)$$

This assumption is also known as *Persistent Excitation* (PE), see e.g. Söderström and Stoica (1989). Note that our assumption in Eq. (4) covers a wider range of cases. For example, Eq. (4) does not require the singular values of $\frac{1}{\sqrt{N}} \mathbf{A}$ to converge, while only requires that they lie in $[c_1, c_2]$ when N increases.

Classically, properties of the Least Square Estimate (LSE) under the model given in Eq. (1) are given by the Gauss-Markov theorem. It says that the Best Linear Unbiased Estimation (BLUE) of \mathbf{x}^0 is the LSE under certain assumptions on the noise term. For the Gauss-Markov theorem, please refer to Plackett (1950). However, the normal LSE does not utilize the ‘sparse’ information of \mathbf{x}^0 , which raises the question that whether it is possible to improve on the normal LSE by exploiting this information. In the literature, several approaches have been suggested, which can perform as if the true support set of \mathbf{x}^0 were known. Such property is termed as the ORACLE property in Fan and Li (2001). In Fan and Li (2001), the SCAD (Smoothly Clipped Absolute Deviation) estimator is presented, which turns out to solve a non-convex optimization problem; later in Zou (2006), the ADALASSO (Adaptive Least Absolute Shrinkage and Selection Operator) estimator is presented. The ADALASSO estimator consists of two steps, which implements a normal LSE in the first step, and then solves a reweighed Lasso optimization problem, which is convex. Recently, in Rojas and Hjalmarsson (2011), two LASSO-based estimators, namely the ‘A-SPARSEVA-AIC-RE’ method and the ‘A-SPARSEVA-BIC-RE’ method, are suggested. Both methods need to do the LSE in the first step, then solve a Lasso optimization problem, and finally redo the LSE estimation.

Remark 1 This note concerns the case that \mathbf{x}^0 is a fixed sparse vector. However, when sparse estimators are applied to estimate non-sparse vectors, erratic phenomena could happen. For details, please see the discussions in Leeb and Pötscher (2008); Kale (1985).

In this note, we will present another approach to estimate the sparse vector \mathbf{x}^0 , which also possesses the ORACLE property with a lighter computational cost. The proposed method consists of three steps, in the first step, a normal LSE is conducted, the second step is to solve a LP (Linear Programming) problem, whose solution is given by a soft-thresholding step, finally, redo the LSE based on the support set of the estimated vector from the previous LP problem. Details will be given in Section 2.

In the following, the lower bold case will be used to denote a vector and capital bold characters are used to denote matrices. The subsequent sections are organized as follows. In section 2, we will describe the algorithm in detail and an analytical solution to the LP problem is given. In Section 3, we will analyze the algorithm in detail. In Section 4, we conduct several examples to illustrate the efficacy of the proposed algorithm and compare the proposed algorithm with other algorithms. Finally, we draw conclusions of the note.

2 Algorithm Description

The algorithm consists of the following three steps:

- *LSE*: Compute the LSE of \mathbf{x}^0 , denoted as \mathbf{x}^{ls} .
- *LP*: Choose $0 < \epsilon < 1$ and solve the following Linear Programming problem:

$$\mathbf{x}^{lp} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\mathbf{x} - \mathbf{x}^{ls}\|_\infty \leq \lambda, \quad (6)$$

where $\lambda = \sqrt{\frac{2n}{N^{1-\epsilon}}}$. Detect the support set \mathcal{T}^{lp} of \mathbf{x}^{lp} .

- *RE-LSE*: Compute the LSE of \mathbf{x}^0 based on \mathcal{T}^{lp} . Form the matrix $\mathbf{A}_{\mathcal{T}^{lp}}$, which contains the columns of \mathbf{A} indexed by \mathcal{T}^{lp} and let $\mathbf{A}_{\mathcal{T}^{lp}}^\dagger$ denote its pseudo-inverse. Then the final estimation \mathbf{x}^{rels} is given by $\mathbf{x}_{\mathcal{T}^{lp}}^{rels} = \mathbf{A}_{\mathcal{T}^{lp}}^\dagger \mathbf{y}$, and $\mathbf{x}_{\mathcal{T}^{lp}^C}^{rels} = \mathbf{0}$, in which \mathcal{T}^{lp}^C denotes the complement set of \mathcal{T}^{lp} .

Note that the LP problem has an analytical solution. Writing the ∞ norm constraint explicitly as

$$\begin{aligned} \mathbf{x}^{lp} = \arg \min_{\mathbf{x}} \sum_{i=1}^n |x_i| \\ \text{s.t. } |x_i - x_i^{ls}| \leq \lambda, \text{ for } i = 1 \dots n. \end{aligned} \quad (7)$$

We can see that there are no cross terms in both the objective function and the constraint inequalities, so each component can be optimized separately. From this observation, the solution of the LP problem is given as

$$x_i^{lp} = \begin{cases} 0, & \text{if } |x_i^{ls}| \leq \lambda \\ x_i^{ls} - \lambda, & \text{if } x_i^{ls} > \lambda \\ x_i^{ls} + \lambda, & \text{if } x_i^{ls} < -\lambda \end{cases}$$

for $i = 1, 2, \dots, n$. Such a solution \mathbf{x}^{lp} is also referred to as an application of the soft-thresholding operation to \mathbf{x}^{ls} , see e.g. Donoho and Johnstone (1995). Several remarks related to the algorithm are given as follows.

Remark 2 Note that the tuning parameter λ chosen as $\lambda^2 = \frac{2n}{N^{1-\epsilon}}$ is very similar to the one (which is proportional to $\frac{2n}{N}$) as given in Rojas and Hjalmarsson (2011) based on the Akaike's Information Criterion (AIC).

Remark 3 The order of λ chosen as $-\frac{1}{2} + \frac{\epsilon}{2}$ is essential to make the asymptotical oracle property hold. Intuitively speaking, such a choice can make the following two facts hold.

- (1) Whenever $\epsilon > 0$, \mathbf{x}^0 will lie in the feasible region of Eq. (6) with high probability.
- (2) The threshold decreases 'slower' (in the order of N) than the variance of the pseudo noise term $\mathbf{V}\Sigma^{-1}\mathbf{U}^T\mathbf{v}$. With such a choice, it is possible to get a good approximation of the support set of \mathbf{x}^0 in the second step.

Remark 4 Though the formulation of Eq. (6) is inspired by the Dantzig selector in Candés and Tao (2007), there are some differences between them.

- (1) As pointed out by one of the reviewer, both the proposed method and the Dantzig selector lie in the following class

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\mathbf{W}(\mathbf{x} - \mathbf{x}^{ls})\|_\infty \leq \lambda. \quad (8)$$

If \mathbf{W} is chosen as the identity matrix, we obtain the proposed method; If \mathbf{W} is chosen as $\mathbf{A}^T\mathbf{A}$, then we obtain the same formulation as given by the Dantzig selector.

- (2) As pointed out in Efron (2007), the solution path of the Dantzig selector behaves erratically with respect to the value of the regularization parameter. However, the solution path of Eq. (6) with respect to the value of λ behaves regularly, which is due to the fact that, given λ , the solution to Eq. (6) is given by the application of the soft-thresholding operation to the LSE estimation. When λ increases, the solution will decrease (or increase) linearly and when it hits zero, it will remain to be zero. This in turn implies computational advances when trying to find a s -sparse solution for given s . A simple illustration of the solution path is given. Assume that $n = 4$ and $\mathbf{x}^{ls} = [2, 0.5, -1, -1.5]^T$, then the solution path to Eq. (6) w.r.t. λ is given in Fig. 1).

Remark 5 From a computational point of view, the SCAD method needs to solve a non-convex optimization problem which will suffer from the multiple local minima, see the discussions in Trevor, Hastie, Tibshirani and Friedman




Fig. 1. An illustration of the solution path to Eq. (6) w.r.t. λ . When λ equals zero, the solution to Eq. (6) is \mathbf{x}^{ls} ; when λ increases, the solution trajectory shrinks linearly to zero and then remains zero.

Table 1
Computational steps needed for different methods

	Step 1	Step 2	Step 3
LP + Re-LSE	LSE	ST	Re-LSE
ADALASSO	LSE	LASSO	
A-SPARSEVA-AIC-RE	LSE	LASSO	Re-LSE
A-SPARSEVA-BIC-RE	LSE	LASSO	Re-LSE

(2005). Hence, the proposed scheme is mainly compared with techniques which can be solved as convex optimization problems. In Table 1, we list the computational steps needed for different methods. In the table, the term ST means the soft-thresholding operation, the term Re-LSE means 'redo the LSE estimation after detecting the support set of the result obtained from the second step'. For a more precise description, see the Algorithm Description section. From this table, we can see that in the first step, all the methods need to do a LSE estimation; in the second step, except the proposed method (which is denoted by LP + Re-LSE), the other methods need to solve a LASSO optimization problem, which is more computationally involved than a simple soft-thresholding operation as needed by the proposed method; except the ADALASSO method, the other methods need to do a Re-LSE step. From this table, we can also see that the main computational burden for the proposed method comes from the LSE step.

Remark 6 Note that the proposed method does not need an "adaptive step" (i.e. to reweigh the cost function) in order to achieve the ORACLE property, which is different from the methods presented in Rojas and Hjalmarsson (2011) and Zou (2006).

3 Analysis of the algorithm

In this section, we will discuss the properties of the presented estimator. In the following, we will denote the smallest singular value of \mathbf{A} as σ .

Remark 7 *In the following sections, we assume that the noise variance equals one, i.e. $c = 1$, for the following reasons:*

- (1) *When the noise variance is given in advance, one can always re-scale the problem accordingly.*
- (2) *Even if the noise variance is not known explicitly (but is known to be finite), the support of \mathbf{x}^0 will be recovered asymptotically. This is a direct consequence of the fact that finite, constant scalings do not affect the asymptotic statements, i.e. we can use the same λ for any level of variance without influencing the asymptotic behavior.*

The following facts (Lemma 1-3) will be needed for subsequent analysis. Since their proofs are standard, we state them without proofs here. Using the notations as introduced before, one has that

Lemma 1 $\mathbf{x}^{ls} = \mathbf{x}^0 + \mathbf{V}\Sigma^{-1}\mathbf{U}^T\mathbf{v}$.

Lemma 2 $\mathbf{b} = \Sigma\mathbf{V}^T\mathbf{x}^{ls} - \Sigma\mathbf{V}^T\mathbf{x}^0$ is a Gaussian random vector with distribution $\mathcal{N}(0, I)$.

Lemma 3 *Given $d > 0$, then*

$$\int_{|t|>d} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt \leq e^{-\frac{d^2}{2}}.$$

In the following, we will first analyze the probability that \mathbf{x}^0 lies in the constraints set of the LP problem given by Eq. (6). Then we give an error estimation of the results given by Eq. (6). After this, we will discuss the capability of recovering the support set of \mathbf{x}^0 by Eq. (6), which will lead to the asymptotic ORACLE property of the proposed estimator.

Lemma 4 *For all $\lambda > 0$, one has that*

$$\mathbb{P}\left(\|\mathbf{V}^T\mathbf{x}^{ls} - \mathbf{V}^T\mathbf{x}^0\|_\infty > \frac{\lambda}{\sqrt{n}}\right) \leq ne^{-\frac{\lambda^2\sigma^2}{2n}}.$$

Proof By Lemma 2, and noticing that $\mathbf{b} = \Sigma\mathbf{V}^T\mathbf{x}^{ls} - \Sigma\mathbf{V}^T\mathbf{x}^0$ is a Gaussian random vector with distribution

$\mathcal{N}(0, I)$, we have that

$$\begin{aligned} & \mathbb{P}\left(\|\mathbf{V}^T\mathbf{x}^{ls} - \mathbf{V}^T\mathbf{x}^0\|_\infty > \frac{\lambda}{\sqrt{n}}\right) \\ & \leq \mathbb{P}\left(\|\Sigma\mathbf{V}^T\mathbf{x}^{ls} - \Sigma\mathbf{V}^T\mathbf{x}^0\|_\infty > \frac{\lambda\sigma}{\sqrt{n}}\right) \\ & = \mathbb{P}\left(\|\mathbf{b}\|_\infty > \frac{\lambda\sigma}{\sqrt{n}}\right) \\ & = \mathbb{P}\left(\exists i, \text{ such that } |b_i| > \frac{\lambda\sigma}{\sqrt{n}}\right) \\ & \leq \sum_{i=1}^n \mathbb{P}\left(|b_i| > \frac{\lambda\sigma}{\sqrt{n}}\right). \end{aligned}$$

Application of Lemma 3 gives the desired result. \square

Lemma 5 *For all $\lambda > 0$, if $\|\mathbf{V}^T\mathbf{x}^{ls} - \mathbf{V}^T\mathbf{x}^0\|_\infty \leq \frac{\lambda}{\sqrt{n}}$, then $\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty \leq \lambda$.*

Proof Define \mathbf{c} as $\mathbf{c} = \mathbf{V}^T\mathbf{x}^{ls} - \mathbf{V}^T\mathbf{x}^0$, so we have $\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty = \|\mathbf{V}\mathbf{c}\|_\infty$. Analyze the i th element of $\mathbf{V}\mathbf{c}$ that

$$|\mathbf{V}_i\mathbf{c}| \leq \|\mathbf{c}\|_2 \leq \|\mathbf{c}\|_\infty \sqrt{n} \leq \lambda.$$

The first inequality is by definition, the second inequality comes from the Cauchy inequality, the last inequality is due to the assumption of the lemma. \square

Combining the previous two lemmas gives

Lemma 6 $\mathbb{P}(\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty \leq \lambda) \geq 1 - ne^{-\frac{\lambda^2\sigma^2}{2n}}.$

Proof The proof goes as follows

$$\begin{aligned} & \mathbb{P}(\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty \leq \lambda) \\ & \geq \mathbb{P}\left(\|\mathbf{V}^T\mathbf{x}^{ls} - \mathbf{V}^T\mathbf{x}^0\|_\infty \leq \frac{\lambda}{\sqrt{n}}\right) \\ & = 1 - \mathbb{P}\left(\|\mathbf{V}^T\mathbf{x}^{ls} - \mathbf{V}^T\mathbf{x}^0\|_\infty > \frac{\lambda}{\sqrt{n}}\right) \\ & \geq 1 - ne^{-\frac{\lambda^2\sigma^2}{2n}} \end{aligned}$$

The first inequality comes from Lemma 5, and the second inequality follows from Lemma 4. \square

The above lemma tells us that \mathbf{x}^0 will lie inside the feasible set of the LP problem as given in Eq. (6) with high probability. By a proper choice of λ , the following result is concluded.

Theorem 1 *Given $0 < \epsilon < 1$, and let $\lambda^2 = \frac{2n}{N^{1-\epsilon}}$, we have that*

$$\mathbb{P}(\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty \leq \lambda) \geq 1 - ne^{-c_1^2 N^\epsilon}.$$

Next, we will derive an error bound (in the l_2 - norm) of the estimator given by the LP formulation. Define

$$\mathbf{h} = \mathbf{x}^{lp} - \mathbf{x}^0,$$

as the error vector of LP formulation. We have that the error term \mathbf{h} is bounded as follows:

Lemma 7 *For any $\lambda > 0$, if $\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty \leq \lambda$, then we have that $\|\mathbf{h}\|_2^2 \leq 4s\lambda^2$.*

Proof We first consider the error vector on \mathcal{T}^c which is given by $\mathbf{h}_{\mathcal{T}^c}$. Since $\|\mathbf{x}^{ls} - \mathbf{x}^0\|_\infty \leq \lambda$ and $\mathbf{x}_{\mathcal{T}^c}^0 = \mathbf{0}$, we have that $\|\mathbf{x}_{\mathcal{T}^c}^{ls}\|_\infty \leq \lambda$. It follows from the previous discussions that \mathbf{x}^{lp} is obtained by application of the soft-shresholding operator with the threshold λ , applied componentwise to \mathbf{x}^{ls} , hence we obtain that $\mathbf{x}_{\mathcal{T}^c}^{lp} = \mathbf{0}$. This implies that $\mathbf{h}_{\mathcal{T}^c} = \mathbf{0}$.

Next we consider the error vector on the support \mathcal{T} , denoted as $\mathbf{h}_{\mathcal{T}}$. From the property of the soft-thresholding operation, it follows that $\|\mathbf{x}_{\mathcal{T}}^{ls} - \mathbf{x}_{\mathcal{T}}^{lp}\|_\infty \leq \lambda$. Then we have that $\|\mathbf{x}_{\mathcal{T}}^0 - \mathbf{x}_{\mathcal{T}}^{lp}\|_\infty \leq \|\mathbf{x}_{\mathcal{T}}^{ls} - \mathbf{x}_{\mathcal{T}}^{lp}\|_\infty + \|\mathbf{x}_{\mathcal{T}}^{ls} - \mathbf{x}_{\mathcal{T}}^0\|_\infty \leq 2\lambda$

Combining both statements gives that $\|\mathbf{h}\|_2^2 = \|\mathbf{h}_{\mathcal{T}}\|_2^2 + \|\mathbf{h}_{\mathcal{T}^c}\|_2^2 \leq |T|\|\mathbf{h}_{\mathcal{T}}\|_\infty^2 \leq 4s\lambda^2$. \square

Plugging in the λ as chosen in previous section, we can get the error bound of the LP formulation. However, the estimate \mathbf{x}^{lp} is not the final estimation, instead it will be used to recover the support set of \mathbf{x}^0 . The following theorem states this result formally. For notational convenience, $\mathcal{T}^{lp}(N)$ is used to denote the recovered support from the LP formulation, and $\mathbf{x}^{rels}(N)$ then denotes the estimate after the second LSE step using N observations. Finally, the vector $\mathbf{x}^{ls-or}(N)$ denotes the LSE as if the support of \mathbf{x}^0 were known (i.e. the ORACLE presents) using N observations.

We will first get a weak support recovery result and based on this, we further prove that the support as recovered by the LP formulation will converge to the true support \mathcal{T} with probability 1 when N goes to infinity.

Lemma 8 *Given $0 < \epsilon < 1$, and assume that the matrix \mathbf{A} has singular values which satisfies Eq. (4), with constants c_1, c_2 as given there. Let $x_0 \triangleq \min\{|x_i^0|, i \in \mathcal{T}\} \in \mathbb{R}^+$, and $\lambda^2 = \frac{2n}{N^{1-\epsilon}}$, then*

$$\lim_{N \rightarrow \infty} \mathbb{P}(\mathcal{T} = \mathcal{T}^{lp}(N)) = 1.$$

Proof Let the vector $\bar{\mathbf{v}}$ denote $\bar{\mathbf{v}} = \mathbf{V}\Sigma^{-1}\mathbf{U}^T\mathbf{v}$. Since $\mathbf{x}^{ls} = \mathbf{x}^0 + \mathbf{V}\Sigma^{-1}\mathbf{U}^T\mathbf{v}$, one has that $\mathbf{x}^{ls} = \mathbf{x}^0 + \bar{\mathbf{v}}$, in which $\bar{\mathbf{v}}$ follows a normal distribution $\mathcal{N}(0, \mathbf{V}\Sigma^{-2}\mathbf{V}^T)$.

Without loss of generality, assume that $x_1^0, x_2^0, \dots, x_s^0$ are the nonzero elements of \mathbf{x}^0 and their values are positive. Since λ decreases when N increases, so there exist a number $N_1 \in \mathbb{N}$, such that $\lambda < \frac{x_0}{2}$ for all $N \geq N_1$. In the following derivations, we use $v_{i,j}$ to denote the element in the i th row, j th column of \mathbf{V} and \bar{v}_i denotes the i th element of $\bar{\mathbf{v}}$. When $N > N_1$, we have the following bound of $\mathbb{P}(\mathcal{T} \neq \mathcal{T}^{lp}(N))$:

$$\begin{aligned} & \mathbb{P}(\mathcal{T} \neq \mathcal{T}^{lp}(N)) \\ &= \mathbb{P}(|x_1^0 + \bar{v}_1| < \lambda, \text{ or } |x_2^0 + \bar{v}_2| < \lambda, \dots, \text{ or } |x_s^0 + \bar{v}_s| < \lambda; \\ & \quad \text{or } |\bar{v}_{s+1}| > \lambda, \text{ or } |\bar{v}_{s+2}| > \lambda, \dots, \text{ or } |\bar{v}_N| > \lambda) \\ &\leq \sum_{i=1}^s \mathbb{P}(-\lambda - x_i^0 < \bar{v}_i < \lambda - x_i^0) + \sum_{i=s+1}^N \mathbb{P}(|\bar{v}_i| > \lambda) \\ &\leq \sum_{i=1}^s \frac{2\lambda e^{-(2\sum_{j=1}^n \sigma_j^{-2} v_{ij}^2)^{-1}(-x_i^0 + \lambda)^2}}{\sqrt{2\pi(\sum_{j=1}^n \sigma_j^{-2} v_{ij}^2)}} \\ & \quad + \sum_{i=s+1}^N e^{-(2\sum_{j=1}^n \sigma_j^{-2} v_{ij}^2)^{-1}\lambda^2} \\ &\leq \sum_{i=1}^s \frac{2c_2\sqrt{N}\lambda}{\sqrt{2\pi}} e^{-\frac{1}{2}c_1^2 N(-x_i^0 + \lambda)^2} + \sum_{i=s+1}^N e^{-\frac{1}{2}c_1^2 N\lambda^2} \\ &\leq 2c_2s\sqrt{n}N^{\frac{\epsilon}{2}} e^{-\frac{1}{8}(c_1x_0)^2N} + Ne^{-c_1^2 nN^\epsilon} \\ &= CN^{\frac{\epsilon}{2}} e^{-\frac{1}{8}(c_1x_0)^2N} + Ne^{-c_1^2 nN^\epsilon}, \end{aligned} \tag{9}$$

where $C = 2c_2s\sqrt{n}$. The second inequality in the chain holds due to the fact that the probability distribution function of \bar{v}_i is monotonically increasing in the interval $[-\lambda - x_i^0, \lambda - x_i^0]$, together with results in Lemma 3.

Then we can see that both terms in (9) will tend to 0 as $N \rightarrow \infty$ for any fixed $\epsilon > 0$, i.e. $\lim_{N \rightarrow \infty} \mathbb{P}(\mathcal{T}^{lp}(N) = \mathcal{T}) = 1$. \square

Remark 8 *Notice the fact that*

$$\mathbb{P}(\mathbf{x}^{rels}(N) = \mathbf{x}^{ls-or}(N)) \geq \mathbb{P}(\mathcal{T}^{lp}(N) = \mathcal{T}),$$

and from the previous Lemma, we know that the right hand side will tend to 1 as N tends to infinity, so it also holds that

$$\lim_{N \rightarrow \infty} \mathbb{P}(\mathbf{x}^{rels}(N) = \mathbf{x}^{ls-or}(N)) = 1.$$

Based on the previous lemma, we have

Theorem 2 *Given $0 < \epsilon < 1$, and assume that the matrix \mathbf{A} has singular values which satisfies Eq. (4), with constants c_1, c_2 as given there. Let $x_0 \triangleq \min\{|x_i^0|, i \in \mathcal{T}\} \in \mathbb{R}^+$, and $\lambda^2 = \frac{2n}{N^{1-\epsilon}}$, then it holds that*

$$\mathbb{P}(\exists N' \text{ such that } \cap_{N=N'}^\infty \{\mathcal{T}^{lp}(N) = \mathcal{T}\}) = 1.$$

Proof From the proof in the previous lemma, we have that when $N > N_1$

$$\begin{aligned} \mathbb{P}(\mathcal{T} \neq \mathcal{T}^{lp}(N)) &\leq CN^{\frac{\epsilon}{2}} e^{-\frac{1}{8}(c_1 x_0)^2 N} + Ne^{-c_1^2 n N^\epsilon} \\ &= Ce^{-\frac{1}{8}(c_1 x_0)^2 N + \frac{\epsilon}{2} \ln(N)} + e^{\ln(N) - c_1^2 n N^\epsilon} \\ &= Ce^{(c_1 x_0)^2 N (\frac{\epsilon \ln(N)}{2(c_1 x_0)^2 N} - \frac{1}{8})} + e^{c_1^2 n N^\epsilon (\frac{\ln(N)}{c_1^2 n N^\epsilon} - 1)}. \end{aligned}$$

Since $0 < \epsilon < 1$ and $x_0 > 0$, one has that $\frac{\epsilon \ln(N)}{2(c_1 x_0)^2 N}$ and $\frac{\ln(N)}{c_1^2 n N^\epsilon}$ will tend to zero if $N \rightarrow \infty$. Hence there exists a number $N_2 \in \mathbb{N}$ such that for all $N > N_3 \triangleq \max(N_1, N_2)$ one has that $\frac{\epsilon \ln(N)}{2(c_1 x_0)^2 N} < \frac{1}{16}$ and $\frac{\ln(N)}{c_1^2 n N^\epsilon} < \frac{1}{2}$. Hence

$$\begin{aligned} &\sum_{N=N_3}^{\infty} \mathbb{P}(\mathcal{T}^{lp}(N) \neq \mathcal{T}) \\ &\leq \sum_{N=N_3}^{\infty} Ce^{-\frac{1}{16}(c_1 x_0)^2 N} + \sum_{N=N_3}^{\infty} e^{-\frac{1}{2}c_1^2 n N^\epsilon} \\ &\leq \int_{N=N_3-1}^{\infty} Ce^{-\frac{1}{16}(c_1 x_0)^2 t} dt + \int_{N_3-1}^{\infty} e^{-\frac{1}{2}c_1^2 n t^\epsilon} dt \\ &= A + B. \end{aligned}$$

Furthermore, it can be seen that

$$A = \int_{N=N_3-1}^{\infty} Ce^{-\frac{1}{16}(c_1 x_0)^2 t} dt < \infty.$$

In the following, we will show that $B = \int_{N_3-1}^{\infty} e^{-\frac{1}{2}c_1^2 n t^\epsilon} dt < \infty$. By a change of variable using $x = \frac{1}{2}c_1^2 n t^\epsilon$, we have that

$$B = \frac{1}{c_1^2 n \epsilon} \int_{\frac{1}{2}c_1^2 n (N_3-1)^\epsilon}^{\infty} x^{\frac{1}{\epsilon}-1} e^{-x} dx < \frac{1}{c_1^2 n \epsilon} \Gamma\left(\frac{1}{\epsilon}\right) < \infty$$

with Γ the Gamma function. And hence

$$\sum_{N=N_3}^{\infty} \mathbb{P}(\mathcal{T}^{lp}(N) \neq \mathcal{T}) < \infty.$$

Application of the Borel-Cantelli lemma [4] implies that the events in $\{\mathcal{T} \neq \mathcal{T}^{lp}(N)\}_{N=N_3}^{\infty}$ will not happen infinitely often, which concludes the result. \square

4 Illustrative Experiments

This section supports the findings in the previous section with numerical examples and make the comparisons with the other algorithms which possess the ORACLE property in the literature.

4.1 Experiment 1

This example is taken from Zou (2006). The setups are repeated as follows.

- \mathbf{x}^0 is set to be $(3, 1.5, 0, 0, 2, 0, 0, 0)^T$;
- Rows of matrix A are i.i.d. normal vectors;
- The correlation between the j_1 -th and the j_2 -th elements of each row are given as $0.5^{|j_1-j_2|}$;
- The noise term $\mathbf{v} \in \mathbb{R}^N$ follows distribution $\mathcal{N}(0, I_N)$.

Based on these setups, the proposed method and also the methods presented in Rojas and Hjalmarsen (2011) (the A-SPARSEVA-AIC-RE method and the A-SPARSEVA-BIC-RE methods) and Zou (2006) (the ADALASSO method) are applied to recover \mathbf{x}^0 . In this experiment, ϵ for the proposed method is set to $\frac{1}{3}$; λ_N for 'ADALASSO' is chosen as $N^{1/2-\gamma/4}$ (this choice satisfies all the assumptions in Theorem 2 in Zou (2006)), and γ is set to 1; the thresholding value (for detecting zero components from the solution of the Lasso problem) for the 'A-SPARSEVA-AIC-RE' and 'A-SPARSEVA-BIC-RE' are set to be 10^{-5} as suggested in Rojas and Hjalmarsen (2011). For the comparison, we also include the experiment result obtained by using the LASSO method, in which we set the tuning parameter as \sqrt{N} . In Fig. 2, for every N , experiment is repeated 50 times to get the estimated MSE. The following abbreviations are used in Fig. 2: (1) the curve with tag 'LSE' gives the MSE of the estimates by the LSE algorithm; (2) the curve with tag 'LP + RE-LSE' gives the MSE of the estimates given by the proposed algorithm; (3) the curve with tag 'ORACLE-LSE' gives the MSE of the estimates by the ORACLE LSE; (4) the curves with tags 'A-SPARSEVA-AIC-RE' and 'A-SPARSEVA-BIC-RE' give the MSE of the estimates by the methods presented in Rojas and Hjalmarsen (2011); (5) the curve with tag 'ADALASSO' gives the MSE of the estimates by the ADALASSO method presented in Zou (2006); (6) the curve with tag 'LASSO' gives the MSE of the estimates of the LASSO method.

Note that, when N becomes large, the curves 'LP + RE-LSE' and 'ORACLE-LSE' exactly match each other.

Fig. 3 demonstrates the efficacy of support recovery of the LP formulation in Eq. (6) for different choices of ϵ . In the plot, 'portion' is defined as the ratio of successful trials over the total number of trials. We conclude the empirical observations for this experiment in the caption of the figure.

In practice, cross validation technique could be exploited to choose the tuning parameters. In the following, we will take the ADALASSO and the proposed method for granted to illustrate the idea and compare the performances for both methods when the parameters are obtained by the cross validation technique.




Fig. 2. Performance of the different estimators from N observations to estimate \mathbf{x}^0 . This picture indicates that the proposed estimator will give exactly the same performance as the ORACLE estimator for a large N ($N \approx 75$).




Fig. 3. Support recovery performance of Eq. (6) for different choices of ϵ . Empirically, we observe that: 1) When ϵ is chosen to be small, the ratio for successful support recovery will be larger when N is small; but when N is large, the ratio for successful support recovery will converge slower to 100% and oscillation exists. This can be observed in the zoomed-in part. 2) When ϵ is chosen to be large, the ratio for successful support recovery will be smaller when N is small; but when N is large, the ratio for successful support recovery will go faster to 100% and no oscillation exists, see also the zoomed-in part in the figure.

In the ADALASSO algorithm and the proposed algorithm, there are two tuning parameters, namely γ for the ADALASSO, and ϵ for the proposed method. In the following part, we will apply the 5-fold cross-validation method (see Trevor, Hastie, Tibshirani and Friedman (2005)) to choose the tuning parameters and then compare their performances based on the chosen tuning parameters. The procedure is as follows. At first, the tuning parameter is obtained by 5-fold cross validation, then it is applied to an independently generated test data which has the same dimension as the training

data and the evaluation data. For different N , we run 100 i.i.d. realizations. In each realization, we record the value $\|\hat{\mathbf{x}} - \mathbf{x}^0\|_2^2$, where $\hat{\mathbf{x}}$ denotes the estimate obtained by the estimator. ϵ are selected from $\{1/8, 1/4, 1/2\}$, γ are selected from $\{1/2, 1, 2\}$, and N are chosen from $\{20, 50, 100, 200, 300, 500\}$. The results are reported in Fig. 4.

4.2 Experiment 2

In this part, we perform an experiment for recovering the sinusoids from noisy measurements. The data is generated as follows:

$$y(t) = \sum_{k'=1}^{n'} c_{i_{k'}} \sin(w_{i_{k'}} t) + v(t).$$

Here both $\{w_{i_{k'}}\}_{k'}$ and $\{c_{i_{k'}}\}_{k'}$ are unknown, but we know that the frequencies do belong to a (larger, but of constant size) set $\{w_k\}_{k=1}^n$ of n elements. By sampling the system with period t_s , we obtain the system

$$\mathbf{y} = \mathbf{A}\mathbf{c}^0 + \mathbf{v}, \quad (10)$$

where $\mathbf{y} = [y(t_s), \dots, y(Nt_s)]^T$. The matrix $\mathbf{A} \in \mathbb{R}^{N \times n}$ is defined as follows. The i -th row of \mathbf{A} is given by

$$\mathbf{A}_i = [\sin(iw_1 t_s), \sin(iw_2 t_s), \dots, \sin(iw_n t_s)], \quad (11)$$

for $i = 1, \dots, N$. The parameter term and noise term are defined as $\mathbf{c}^0 = [c_1, c_2, \dots, c_n]^T$, and $\mathbf{v} = [v(t_s), v(2t_s), \dots, v(Nt_s)]^T$.

In this experiment, $n = 10$ and $\mathbf{c}^0 = (1, 1, 1, 0, \dots, 0)^T$, $w_k = k$ for $k = 1, 2, \dots, n$. We increase N up to 500 and the noise vector \mathbf{v} satisfies $\mathbf{v} \sim \mathcal{N}(0, I_N)$. We also assume that only the first three entries in $\{w_k\}_{k=1}^n$ occur effectively in the system of Eq. (10) and the corresponding amplitudes are set to 1, i.e. $n' = 3$ and $i_1 = 1$, $i_2 = 2$, $i_3 = 3$. The sampling period t_s is set to $0.1s$.

The result using the proposed algorithm to recover \mathbf{x}^0 is displayed in Fig. 5. It is again clear that the proposed estimator is as efficient as the ORACLE estimator if one has enough samples.

This is indeed predicted by the theory above since the \mathbf{A} in Eq. (10) obeys the assumption of Eq. (4). This follows from the proposition given as:

Proposition 1 *There exist constants $\{C_{i,j}\}_{0 \leq i,j \leq n}$ which do not depend on N , such that the following results hold. For any $1 \leq i \neq j \leq n$, one has that:*

$$|(A^T A)_{i,j}| = \left| \sum_{t=1}^N \sin(tw_i t_s) \sin(tw_j t_s) \right| \leq C_{i,j} \quad (12)$$

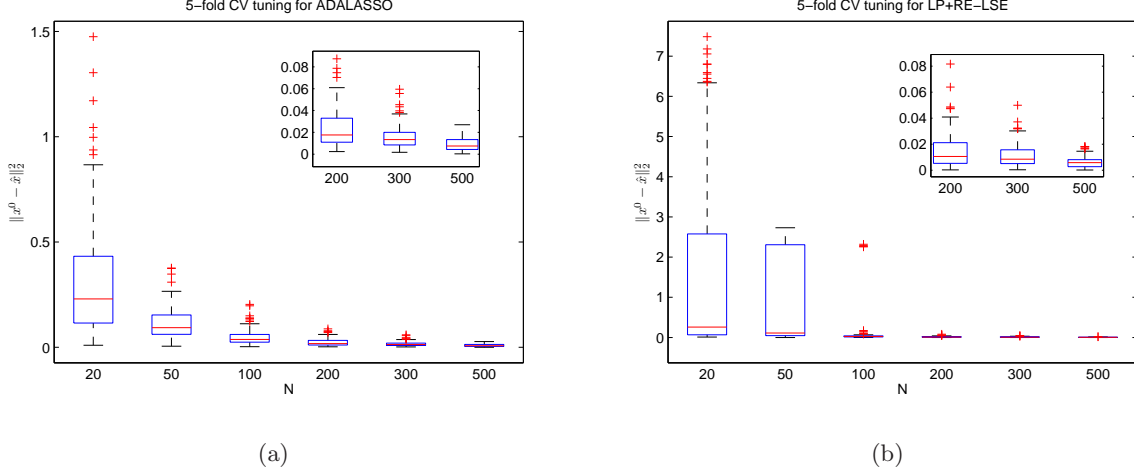


Fig. 4. This figure demonstrates the boxplots of the recovery error obtained through the ADALASSO estimator and the proposed estimator when the tuning parameters are chosen by the 5-fold cross validation method. From this figure, we can see that performances of both methods are similar when N is large, see the zoomed-in part in the figures. It can also be observed that when N is small, the ADALASSO method has smaller recovery error compared with the proposed method.




Fig. 5. Performance of applying the proposed estimator to recovery sinusoids functions from N observations in Experiment 2. This example also indicates that after a finite number the estimate is exactly equal to the ORACLE estimator.

and for any $1 \leq i \leq n$ that:

$$(A^T A)_{i,i} = \sum_{t=1}^N (\sin(tw_i t_s))^2 \geq \frac{N}{2} - C_{i,i}. \quad (13)$$

The proof is given in Appendix A. With this proposition, an application of Geršgorin circle theorem implies that the eigenvalues of $\mathbf{A}^T \mathbf{A}$ will increase with the order of N , which in turn implies Eq. (4).

5 Conclusion

This note presents an algorithm for solving an over-determined linear system from noisy observations, spe-

cializing to the case where the true 'parameter' vector is sparse. The proposed method does not need one to solve explicitly an optimization problem: it rather requires one to compute twice the LSE step, as well to perform a computationally cheap soft-thresholding step. Also, it is shown formally that the proposed method achieves the ORACLE property. An open question is to quantify how many samples would be sufficient to guarantee exact recovery of \mathbf{x}^0 for given sparsity level s . In this note, we resort to the asymptotic Borel-Cantelli Lemma ('there exists such a number'), but it is often of interest to have an explicit characterization of this number. Another open question is how to find a suitable weighting matrix \mathbf{W} which can further improve the performance of the proposed algorithm.

References

- Rojas, C. & Hjalmarsson, H. (2011). Sparse estimation based on a validation criterion. *The 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC11)*, Orlando, USA.
- Donoho, D. L. (2006). Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4), 1289-1306. Germany: De Gruyter.
- Plackett, R. L. (1950). Some theorems in least squares. *Biometrika*, 37(1/2), 149-157.
- Rick, D. (2005). *Probability: Theory and Examples (2nd Edition)*. Duxbury press.
- Donoho, D. L. & Johnstone, I. M. (1995). Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association*, 90(432), 1200-1224.
- Söderström, T. & Stoica, P. (1989). *System Identification*. UK: Prentice-Hall International.
- Fan, J. & Li, R. (2001). Variable selection via nonconcave

penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456), 1348-1360.

Candès, E. & Tao, T. (2007). The Dantzig selector: statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6), 2313-2351.

Kump, P., Bai, Er-W., Chan, K. S., Eichinger, B. & Li K. (2012). Variable selection via RIVAL (removing irrelevant variables amidst Lasso iterations) and its application to nuclear material detection. *Automatica*, 48(9), 2107-2115.

Kukreja, S. L. (2009). Application of a least absolute shrinkage and selection operator to aeroelastic flight test data. *International Journal of Control*, 82(12), 2284-2292.

Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 73(3), 267-288.

Zou, H. (2006). The adaptive Lasso and its oracle properties. *Journal of the American statistical association*, 101(476), 1418-1429.

Efron, B., Hastie, T. & Tibshirani R. (2007). Discussion of 'the Dantzig selector'. *The Annals of Statistics*, 35(6), 2358-2364.

Stoica, P. & Moses, R. L. (1997). *Introduction to spectral analysis (Vol. 89)*. New Jersey: Prentice hall.

Kale, B. K. (1985). A note on the super efficient estimator. *Journal of Statistical Planning and Inference*, 12, 259-263.

Leeb, H. & Pötscher, B. M. (2008). Sparse estimators and the oracle property, or the return of Hodges's estimator. *Journal of Econometrics*, 142(1), 201-211.

Candès, E. J. & Wakin, M. B. (2008). An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2), 21-30.

Trevor, J., Hastie, T., Tibshirani, R. J. & Friedman, J. H. (2005). *The elements of statistical learning: data mining, inference, and prediction*. Springer.

A Proof of Proposition 4

Proof The proof of (12) goes as follows. First

$$\begin{aligned} & \left| \sum_{t=1}^N \sin(tw_i t_s) \sin(tw_j t_s) \right| \\ &= \frac{1}{2} \left| \sum_{t=1}^N (\cos(t(w_i - w_j)t_s) - \cos(t(w_i + w_j)t_s)) \right| \\ &\leq \frac{1}{2} \left| \sum_{t=1}^N \cos(t(w_i - w_j)t_s) \right| + \frac{1}{2} \left| \sum_{t=1}^N \cos(t(w_i + w_j)t_s) \right|. \end{aligned}$$

We focus on bounding the term $\left| \sum_{t=1}^N \cos(t(w_i - w_j)t_s) \right|$, the bound of the other term will follow along the same lines.

$$\begin{aligned} & \left| \sum_{t=1}^N \cos(t(w_i - w_j)t_s) \right| \\ &= \left| \operatorname{Re} \left(\frac{1 - e^{j(N+1)(w_i - w_j)t_s}}{1 - e^{j(w_i - w_j)t_s}} \right) - 1 \right| \\ &\leq \left| \frac{1 - e^{j(N+1)(w_i - w_j)t_s}}{1 - e^{j(w_i - w_j)t_s}} \right| + 1 \\ &\leq \frac{2}{|1 - e^{j(w_i - w_j)t_s}|} + 1, \end{aligned}$$

which is a constant which does not depend on N , so inequality (12) is obtained.

In order to prove inequality (13), observe that

$$\begin{aligned} \sum_{t=1}^N (\sin(tw_i t_s))^2 &= \frac{1}{2} \sum_{t=1}^N (1 - \cos(2tw_i t_s)) \\ &\geq \frac{N}{2} - \frac{1}{2} \left| \sum_{t=1}^N \cos(2tw_i t_s) \right|. \end{aligned}$$

Using previous bounding method, $\frac{1}{2} \left| \sum_{t=1}^N \cos(2tw_i t_s) \right|$ is also bounded by a constant $C_{i,i}$ which does not depend on N . This concludes the proof. \square